N90-25575

EVALUATION OF OFF-ROAD TERRAIN WITH STATIC STEREO AND MONOSCOPIC DISPLAYS

John P. Yorchak
Craig S. Hartley

Martin Marietta Astronautics Group
P.O. Box 179, MS H4371
Denver, Colorado 80201
(303) 971-5748

ABSTRACT

The National Aeronautics and Space Administration is currently funding research into the design of a Mars "rover" vehicle. This unmanned rover will be used to explore a number of scientific and geologic sites on the Martian surface. Since the rover can not be driven from Earth in real-time, due to lengthy communication time delays, a locomotion strategy that optimizes vehicle range and minimizes potential risk must be developed. In order to assess the degree of on-board artificial intelligence (AI) required for a rover to carry out its' mission, we have conducted an experiment to define a "no AI" baseline. In our experiment 24 subjects, divided into stereo and monoscopic groups, were shown video snapshots of four terrain scenes. The subjects' task was to choose a suitable path for the vehicle through each of the four scenes. Paths were scored based on distance travelled and hazard avoidance. Study results are presented with respect to; (1) risk versus range, (2) stereo versus monocular video, (3) vehicle camera height, and (4) camera field-of-view.

INTRODUCTION

The success of the Viking landers on Mars is well documented, but these missions served to raise as many new questions about the Martian surface as they answered. To attempt to answer some of these questions, a broader, more comprehensive exploration of the Martian surface has been proposed. A mission such as this will collect samples from a number of scientific and geologic sites using a "roving" vehicle. This vehicle, capable of some autonomous operation, will be required to navigate the rough and unpredictable terrain of Mars.

A Mars rover can not be driven in real-time from Earth due to communication time delays ranging anywhere from nine to forty minutes, the limited availability of windows for data transmissions, and the limitations of data rates for distances such as this. Therefore, a locomotion strategy that optimizes vehicle range and minimizes potential risk to the vehicle must be developed.

When we speak of the range of the vehicle we must consider several complex issues. If the vehicle requires repeated commands from Earth to travel between sampling sites, it's total range on the surface will be reduced. If on the other hand, the vehicle is capable of autonomous moves, due to onboard artificial intelligence (AI) and a superior ranging and vision system, then the vehicle's range and it's ability to gather more samples, is maximized. Unfortunately, the cost of this autonomous capability is extremely high.

In order to begin to assess how much AI the rover needs onboard to carry out its' mission, we have designed a study to define a "no AI" baseline. This study will evaluate a Mars rover scenario where the vehicle possesses no onboard intelligence, and therefore, would be teleoperated from Earth. With this scenario, the only visual information available to the operator concerning the Martian surface would be a series of video snapshots from the rover's cameras. The key question we have addressed with this experiment is, using only video snapshots, is it possible for an operator to plan a safe path for the vehicle through hazardous terrain, while at the same time, maximizing each discrete vehicle move?

If we look at two operational extremes this question should become a little more clear. On one hand, as an operator of the vehicle I am fairly certain that I can steer clear of hazardous terrain if I move the vehicle one meter at a time. However, since I will only be able to make one move approximately every 40 minutes, the vehicles total range per day would only be 24 meters. If, on the other hand, I take some risks, I might be able to move the vehicle 20 meters at a time. With this scenario I would be able to move 480 meters per day, but the element of risk may be unacceptable.

Using video snapshots of terrain around the main plant at Martin Marietta Astronautics in Denver, Colorado, we have created four scenes that contain differing degrees of perceived hazards and actual hazards. Using these four scenes as stimulus displays, subjects were asked to draw the safest and most direct path through the scene that the vehicle could take. Stereo and monocular snapshots were presented to equal numbers of subjects.

METHOD

Subjects

Two groups of twelve subjects completed the experiment. The subjects were all employees of the Martin Marietta Astronautics Group who volunteered for the experiment. The subjects' positions ranged from upper management to lab technician.

Apparatus

Control Console

This experiment was conducted using one of the Space Operations Simulator (SOS) Laboratory Advanced Technology Control Consoles. These consoles were designed as testbeds for such advanced human/system interface technologies as touchscreens, programmable display pushbuttons, speech recognition, speech synthesis, expert system workaids, hypermedia systems, integration of computer graphics and real-time video imagery, and stereoscopic video displays. As testbeds, these consoles are routinely used to evaluate a wide range of technological innovations in human interface design in the context of real-time task simulations conducted by the SOS Lab.

Structure. The control console configuration used for this experiment is shown in Figure 1 below. A custom design metal structure encloses all of the console components. The structure is optimized to match the anthropometric requirements for 50th percentile male operators. However, the console design can accommodate users ranging from 5th percentile females through 95th percentile males through the provision of a special adjustable chair, adjustable handcontrollers, arm rests, and foot pedals. All controls and displays are positioned and angled to ensure easy access and clear direct visibility from a reference eyepoint at the operator's seated position. Only the lower center monitor was used for this experiment.

For the experiment described here, subjects were asked to sit at a comfortable viewing distance (typically 16 to 18 inches) from the primary, lower-center monitor. They adjusted their chair height and position so that their eyes were normal (at 90 degrees) to the lower center monitor. Handcontrollers were installed but were not used as part of the experiment.

Lower Center (Stereoscopic Display) Monitor. The lower center control console monitor was used as the display for this experiment. It is a special 16-inch Tektronix high-resolution, fast refresh (120 Hz), fast phosphor, color CRT designed for real-time stereoscopic or monoscopic video display. This monitor is optionally equipped with a touchscreen when the stereo system is not in use, but normal control

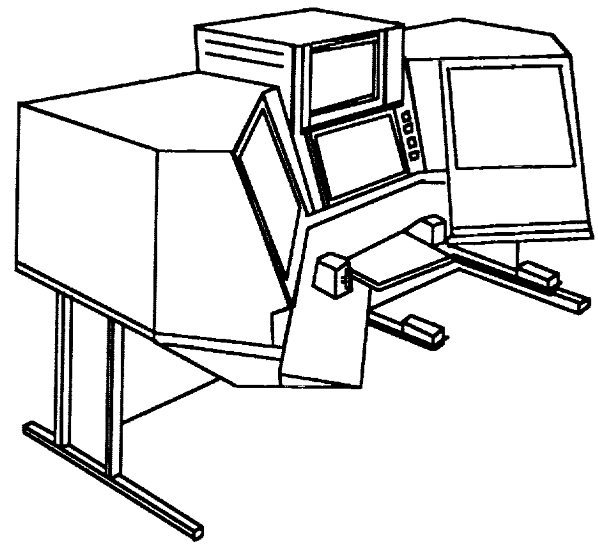

Figure 1. Control console configuration.

of on-screen functions (such as graphics reticle overlay selection) is provided either by programmable LED or EL display pushbuttons located to the right of the monitor or by voice command. The monitor display is driven by Parallax 1280-V-8VN-TS high resolution (1280 x 1024) videographic processor boards mounted in a Sun 3/260 computer.

The stereoscopic video display system uses a time serial presentation of left and right eye views in synch with shifts in polarity of a Tektronix Liquid Crystal Display (LCD) Shutter overlay. The shutter is circularly polarized to allow consistent stereo viewing even when the viewer's head is tilted. A small, horizontal, and nearly invisible seam divides the middle of the screen because it consists of two pieces to facilitate high-speed (120 Hz) switching of the LCD polarity. The shutter is switched by TTL command line from the host computer.

Video images and synchronized TTL switching signals for the LCD shutter are generated by a Parallax Graphics 1280 -V-8VN-TS video graphic board set mounted in a VME bus 21-slot expansion chassis connected to the Sun 3/260 computer. The Parallax video graphics boards provide for input of two NTSC camera views (from the left and right cameras). As the video images are received, the boards digitize each NTSC image into a 640X480 color bit-mapped image (in near real-time) and stores it in a separate RAM buffer. To display stereo, the system alternately sends the bit-mapped image for left and right eye views. At the same time it sends a synchronized TTL command signal to the LCD shutter to switch the polarity as appropriate for each eye view. The operator wears special glasses having each eye lens separately cross polarized so that he/she only sees the left eye view when the digitized image from the left eye is displayed and vice versa. Special software allows shifting of the left and right eye image pixel arrays so that all undesirable vertical and horizontal disparity could be eliminated. This capability allowed us to compensate for slight mechanical misalignments inherent in the camera mounts. The system is capable of maintaining a 15 frame/sec./eye (30 frames/sec. for both eyes) rate with a 640X480 pixel full color image refreshed at 120 Hz. Resolution can be traded for frame rate so it is possible to display a 512 X 240 pixel image at 30 frames/sec./per eye.

The stereo display condition used as an independent variable for this experiment consisted of a static (not moving) 640 X 480 color presentation at 15 frames/sec./eye with a 120Hz refresh rate. The image was previously digitized and stored on hard disk. The image was clear, free of any flicker, and not discernable from live video of the same scene. The stereo effect was excellent and very realistic, and represented the state-of-the-art in stereoscopic display technology.

The monoscopic condition was created by displaying a single (right eye) view from the stereo pair. This image was also flicker free and was identical in appearance to the stereo condition. The only difference in the display was the slight horizontal shift in picture location associated with a single eye view in a stereo pair.

Cameras

The video input devices used for this study consisted of a matched pair of Panasonic CCD high resolution color cameras equipped with Panasonic motor drive auto iris lenses. These systems provided a 57 degree field-of-view when fully zoomed back to their widest angle (the setting used in this experiment).

The cameras were mounted on a custom-designed, computer-controller, mechanism that provided stepper motor control to 0.1 degree in pan, tilt, and camera convergence. The interocular separation between cameras was manually adjustable. Prior to the start of this experiment we carried out a number of pilot studies to establish an optimized stereoscopic picture using this system. Based on that research, literature reviews, consultations with stereo system experts, and the sizes of the system components, we set the interocular distance at 6".

Stereo Camera Setup

Discussions with Dr. Ed Spain at the Naval Oceans System Center (NOSC) and Dr. Robert Cole at the University of Hawaii regarding an optimal stereo camera setup for a remotely controlled vehicle, indicate that an interocular distance of 2.5 - 3", and cameras as close to parallel as possible, is the best setup for distances greater than 100 ft. When the area of interest is 50 - 75 ft. away, a separation of 7.5 - 10" is recommended when the cameras are converged at a point just in front of the "target". For this study, the ALV's stereo camera pair were separated approximately 6", and converged at a point approximately 58 ft. in front of the vehicle. This configuration was used for all four scenes.

As seen in Figure 2, the stereo pair mechanism was mounted on a platform 3 feet above the roof of the ALV, making the camera positions it approximately 15 ft. above the ground. In order to create the widest field of view possible, the cameras were zoomed out all the way (creating a 57 degree horizontal field of view), and the camera platform was tilted down slightly. The bottom of the field-of-view in the console monitor was 24 ft. in front of the vehicle.

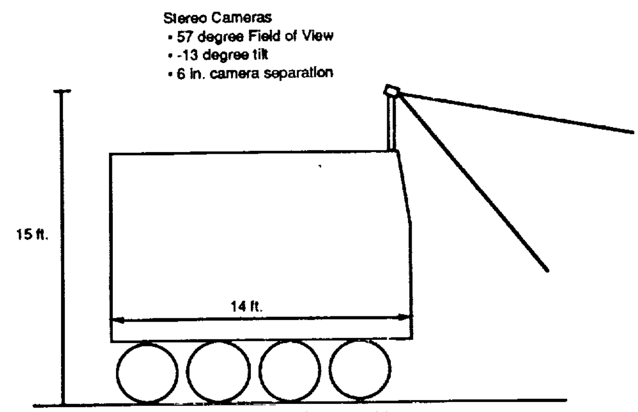

Figure 2. Location and set-up of ALV's stereo camera pair.

Terrain Snapshots

A series of pilot studies were conducted after the stereo camera set up was optimized and prior to the actual experiment to select four test sites that met the following conditions:

Scene 1: actual hazard present in the scene even though it appears relatively safe to the operator,

Scene 2: no actual hazard present in the scene even though the scene appears hazardous,

Scene 3: no actual hazard present in the scene and no appearance of any, and

Scene 4: actual hazard present in the scene and the scene appears hazardous.

Once the sites were selected, video snapshots were made of each by poitioning the vehicle, digitizing the left eye and right eye camera images, and saving those images to disk for later redisplay. All four scenes were digitized on the same day and within a one hour time period so that lighting and weather conditions were the same for each image.

Experimental Procedure

As the subjects arrived at the lab they were greeted by the experimenter and seated at the console so that their eyes were at a standard height and position, with respect to the display (approximately 18" from the screen). They were given the stereo glasses to wear (if they were in the stereo group) and then given a few minutes to view a sample scene in order to get acclimated to the stereo.

As the subjects viewed this scene the experimenter read them the following briefing:

This experiment is being conducted to investigate human remote-driving of the Mars Rover vehicle using stereo (video) snapshots. Since a Mars Rover doesn't currently exist, we are conducting these studies with the functionally similar Autonomous Land Vehicle (ALV). We are attempting to determine the average distance the driver of the Mars Rover would move the vehicle, given only stereo (video) snapshots of the scene. We are also interested in the length, and driver confidence, of these moves given the fact that the scenes often contain unknown hazards.

Your task will be to view 4 different stereo (video) snapshots. After studying the scene for a few moments, you will be asked to draw a path for the vehicle to travel on a line drawing of the scene. This path should be the safest path possible that would take the vehicle from its' current position through the terrain displayed in the snapshot. After drawing the path, place an "X" on the point where you would like the vehicle to make its' first stop to send you a new display. The location of this "X" should represent the point on the path that, given only the information provided by the snapshot, you believe the vehicle could reach without running into any hazards in the terrain.

Are there any questions so far?

Before we start the actual experiment we will do a practice trial on the scene displayed in front of you.

Take a few moments to plan a path through the scene. When you have a path in mind say "OK".

(After the subject says "OK", give him/her a line drawing of the scene to mark on.)

Mark your path on this drawing, and then mark your "X" where you would like the vehicle to stop and send you a new display.

If you don't have any other questions, we're ready to begin the experiment.

Since the subjects had the actual video display in front of them as they used the hardcopy drawing of the scene, they had no trouble drawing their path and positioning their "X" on the picture. We deliberately chose to use a "poor" quality picture because we did not want to add to the subjects' knowledge of the terrain in any way.

In order to be sure that the subjects fully understood their task, the experimenter and the subject worked through the procedure described above using a sample scene displayed on the monitor. This practice scene consisted of a view down a gravel road with hills, telephone poles, and a few bushes and trees in adjacent fields.

An experimental trial consisted of the experimenter sending the appropriate scene to the subjects' monitor from a separate experimenter's workstation, and at the same time starting a timer. After the subject had taken as much time as he/she needed to study the scene and plan a path through it, they told the experimenter that they had a path in mind. The experimenter then stopped the timer and gave the subject the hardcopy picture to draw on. After they returned the marked-up drawing to the experimenter, the next scene was displayed on the subject's monitor. To guard against the possibility of an order effect confounding the data, the presentation order of the scenes was counterbalanced across subjects.

After completing all four experimental trials, the subject completed a brief questionnaire asking about his or her sex, age, position at Martin Marietta, eyesight, experience with stereo vision systems, and experience with remotely operated vehicles (including remotely controlled hobby cars).

After completing the questionnaire each subject was thoroughly debiefed. This debriefing consisted of redisplaying the scenes and soliciting comments about the subjects rationale for their path and placement of their "X". Subjects in the monoscopic group were shown the scenes again, this time in stereo. Many of them remarked that stereo made the scenes look different, but not different enough to change their paths.

Independent Variables

The experimental design used in this study was a mixed factor 2X2X2 repeated measures analysis of variance, where the display viewed, stereo or monocular, was the between subjects factor, and the two within subjects factors were presence of hazard and appearance of hazard. Table I illustrates this experimental design.

Other data, such as sex, age, eyesight, and related experience, gathered from the questionnaire were analyzed with a regression model.

Dependent Variables

Data for three dependent variables were collected and analyzed. These were; (1) the amount of time the subject took to study the scene from the time it was first displayed to when he/she said "OK", (2) the actual distance travelled to the "X", and (3) a driver "efficiency" score derived from a qualitative assessment of the path's direction, distance to the "X", and degree of hazard avoidance.

Data Analysis

Each of the dependent measures was analyzed with a 2X2X2 repeated measures analysis of variance model. The statistical package used to analyze the data was Systat version 3.1 running on a Macintosh II.

	Stereoscopic (12 subjects)		Monoscopic (12 subjects)	
	Appearance of Hazard	No Appearance of Hazard	Appearance of Hazard	No Appearance of Hazard
Hazard Present	Scene 4	Scene 1	Scene 4	Scene 1
No Hazard Present	Scene 2	Scene 3	Scene 2	Scene 3

Table I. Experimental design.

RESULTS

Decision Times

The amount of time subjects took to decide on a path was collected by having the experimenter start a timer once the scene was displayed on the subject's monitor, and then stop it when the subject indicated to the experimenter they had a path.

The results of the analysis of variance indicate that display type did not significantly effect decision times, $F(1,22) = 0.803$, $p<.380$. However, the main effect of appearance of hazards in the scene did, $F(1,22) = 20.259$, $p<.000$. There was also a significant interaction between this main effect and display type, $F(1,22) = 4.383$, $p<.048$. The other main effect of hazard in the scene was not significant, $F(1,22) = 2.595$, $p<.121$. Actual decision times are shown in Table II.

	Appearance of Hazard	No Appearance of Hazard
Hazard Present	Scene 4 (54.2)	Scene 1 (38.5)
No Hazard Present	Scene 2 (51.0)	Scene 3 (29.0)

(Note: Times in parentheses are in seconds.)

Table II. Decision times collapsed across display type.

Distance Traveled to the "X"

The analysis of the distances subjects traveled to the "X" indicate that display type made no difference in their decisions, $F(1,22) = 0.006$, $p<.941$. However, as with decision times, a main effect was observed for the appearance of hazards in the scene, $F(1,22) = 37.815$, $p<.000$. In addition, a main effect for hazards in the scene was also significant, $F(1,22) = 15.273$, $p<.001$. Average distances collapsed across display type are shown in Table III.

	Appearance of Hazard	No Appearance of Hazard
Hazard Present	Scene 4 (22.3)	Scene 1 (32.5)
No Hazard Present	Scene 2 (13.0)	Scene 3 (97.3)

(NOTE: Distances in parentheses are in feet.)

Table III. Distances collapsed across display type.

Efficiency Score

The efficiency score was derived by drawing a grid over a picture of each scene. Each cell of this grid was given a score between 0 and 10 for each of three factors. Factor 1 was the distance driven through the scene, factor 2 was the direction, and factor 3 was the degree of hazard avoidance. Therefore, if a subject placed his/her "X" in the optimal position on a scene, then that cell of the grid was worth 30 points. The analysis conducted on the driver efficiency scores revealed that display type did not significantly effect these scores, $F(1,22) = 0.247$, $p<.624$. A main effect was observed, however, for the appearance of hazards in the scene, $F(1,22) = 63.716$, $p<.000$. The main effect for actual hazards was not significant, but the interaction between appearance and actual hazards was, $F(1,22) = 7.949$, $p<.010$. These scores have been normalized so that a score of 30 represents the best possible score available on each of the four scenes. Therefore, as can be seen in Table IV, the appearance of a hazard in the scene apparently caused subjects to choose paths that were less than optimal.

	Appearance of Hazard	No Appearance of Hazard
Hazard Present	Scene 4 (19.46)	Scene 1 (23.42)
No Hazard Present	Scene 2 (18.21)	Scene 3 (27.67)

Note: The values in parentheses may range from 30 (Best) to 0 (Worst).

Table IV. Efficiency scores collapsed across display type.

Terrain Scenes

A detailed analysis of each scene is required if we are to fully understand the significance of the anovas.

Scene 1

The terrain in this scene is essentially hazard-free in the lower two-thirds of the scene, but the upper third contains an unknown area that most subjects perceived as a potential hazard. Since the majority of the subjects perceived a hazardous area in the scene, it may not have been the best scene for the "No Appearance/Hazard Present" scene. However, five out of our twenty-four subjects (21%) apparently did not see the hazard because they put their "X" in it. Of the 19 who stayed short of the hazard, their average movement was 27 ft., or approximately 75% of the total distance they could have moved. Figure 3 shows the actual distances that subjects moved in this scene.

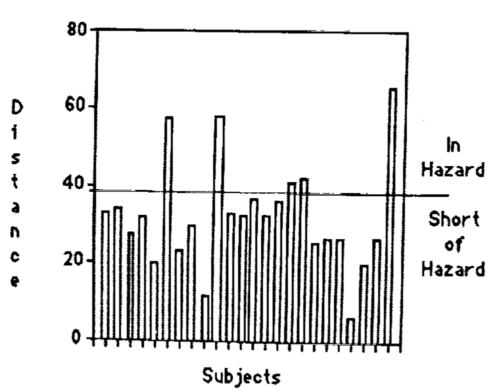

Figure 3. Distance moved by each subject in Scene 1.

Scene 2

Scene 2 was composed to be a hazardous-looking scene that actually contained no hazards. Since there were no actual hazards present in this terrain it is impossible to say whether or not subjects made "risky" decisions regarding their path and distance. However, it is obvious from Figure 4 that most subjects made very conservative decisions regarding this scene. Their difficulty arose from an area directly in front of the vehicle that appeared to be a ditch or a wash. Actually, this area is perfectly safe, with the deepest depression being less than 3" deep. As can be seen in Figure 4, all subjects except one steered around this "problem" and kept their first move quite short. Figure 4 shows the actual distance moved by each subject. The maximum distance possible was 80 feet.

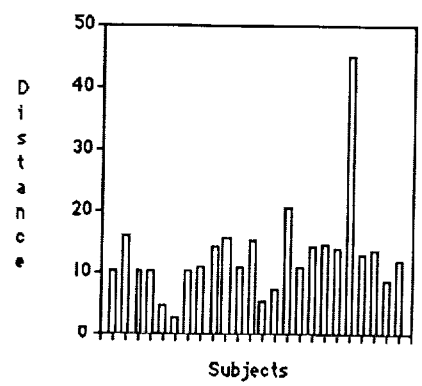

Figure 4. Distance moved by each subject in Scene 2.

Scene 3

Scene 3 was composed over a flat, hazard-free section of the test area, so that it would appear hazard-free to the subjects. It appears from the data that this was the perception of most of the subjects. However, not all subjects perceived the scene as totally hazard-free. As Figure 5 indicates, 13 out of 24 subjects (54%) made conservative judgements regarding this terrain. Figure 5 shows the actual distance moved by each subject. Note that five subjects went as far as they could through the scene.

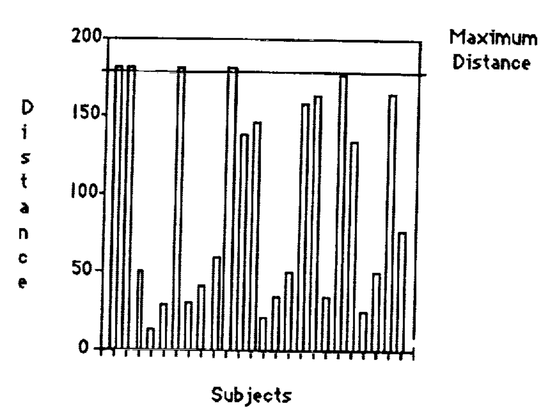

Figure 5. Distance moved by each subject in Scene 3.

Scene 4

Scene 4 was composed to appear hazardous to the subjects while actually being quite hazardous to a vehicle. The hazards present in the scene cut across it at a 30 degree angle to the horizontal. These hazards consisted of a number of channels and washes that would cause a great deal of difficulty for a vehicle. Most subjects perceived the hazards in the upper half of the scene, but they also mis-classified areas directly in front of the vehicle as hazardous when in fact they weren't. Figure 6 shows the actual distance moved by each subject in this scene.

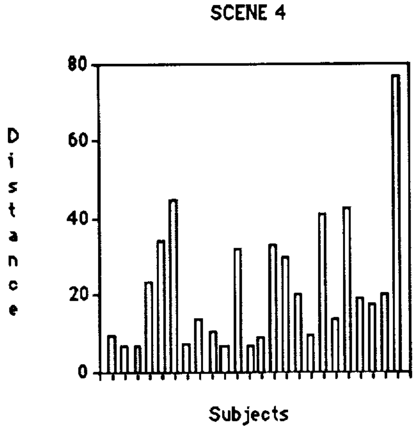

Figure 6. Distance moved by each subject in Scene 4.

DISCUSSION

Risk and Range

Many interpretations of the data presented in this study are possible but we will concentrate on two basic issues---risk and range. With the exception of Scene 1, where five subjects put their "X's" in the hazard, subjects were generally very cautious with their path and distance decisions. It seems that they erred more often in a cautious direction than a careless one. This result is consistent with a finding reported by Spain (1987). In this study he reported that experienced teleoperators of off-road vehicles tended to make fewer errors, but also drove slower. Spain concluded that the effect of training may be learning caution.

Erring on a cautious side would seem to be desirable for operational scenarios involving the Mars rover. However, overly cautious operators significantly decrease vehicle range. In the case where a Mars rover has a limited life on the surface of the planet, it is critical to mission success to maximize each movement. As stated previously, and seen in the scenes that were evaluated, even when hazards were not present in the scene operators often avoided anything that even looked hazardous. This result indicates a need for on-board sensing and ranging of the ground surface. A locomotion strategy that appears reasonable given our findings has been proposed by Wilcox, Salo, Cooper, and Killon (1986). In their study, the operator of a computer-aided remote driving system was asked to identify path points from viewing a 3-dimensional video display of the area in front of the vehicle. Given the path points the operator selected, the computer system commanded the vehicle to execute the path. With this approach the operator is removed from the details of vehicle control, while the computer is removed from the tasks of scene analysis, global path planning, and obstacle avoidance. Testing with this system has indicated that operation over path segments of up to 40 meters has been demonstrated successfully.

Stereo vs. Monocular

It is somewhat surprising that stereo did not aid subjects in the interpretation of terrain scenes. In similar studies, stereo has generally outperformed monocular displays (Spain, 1987). One important difference between this study and those reported by Spain, however, is that our subjects only had single, static frames of stereo or monocular video, rather than real-time dynamic displays. This would seem to indicate that the advantage of stereo lies in the depth cues provided by the relative motion between objects in the foreground and background. When motion is not presented, even though the stereo snapshots "look" better than comparable monocular snapshots, the lack of motion cues on the stereo displays makes them comparable to monocular views.

Discussions with Dr. Robert Cole at the University of Hawaii have revealed another interpretation of our failure to find an advantage for stereo. The height, field of view, and downward tilt of the stereo cameras are critical to the overall stereo effect created. With our stereo pair mounted on top of the ALV approximately 15 ft. above the ground and tilted down 13 degrees, the overall scene was "flattened" to such an extent that the stereo effect was essentially eliminated. This finding is important for future work involving the ALV's stereo system, whereby, in order to optimize the stereo effect provided by the ALV's cameras, it may be necessary to lower the camera platform considerably.

Spain (1987) also concluded that when experienced and inexperienced drivers of a teleoperated vehicle drove with stereo and monocular video, no significant difference was found for stereo, "... one must remember that past research has shown that the advantages stereoscopic imagery provides are most pronounced in unfamiliar, visually cluttered and in visually degraded scenes. Stereoscopic imagery is also useful in judging the relative distances and orientations of objects and terrain surface features - all of which might prove invaluable to an operator in "reading" terrain before attempting to traverse it."

Camera Height

When the stereo pair is too high off the ground, it is impossible to see the horizon and the ground directly in front of the vehicle simultaneously. If the horizon is centered on the display, then a considerable distance in front of the vehicle is not in view. Conversely, if the ground in front of the vehicle is important, the cameras must be tilted down so dramatically that the horizon is not in view and the terrain that is in view becomes "flattened".

These results can easily be applied to Mars rover. Given our finding that stereo is adversely effected when the camera pair is too high off the ground, future studies should be done in order to define the optimum height of the rovers' stereo cameras.

Camera Field of View

Field of view is very important to operators of remote vehicles. As mentioned previously, our cameras were set at the maximum horizontal field of view possible, 57 degrees. This field of view seemed adequate but the height and downward tilt of the ALV's cameras had the effect of causing 22 ft. directly in front of the vehicle to be out of view of the operator. An interaction of camera field of view, height, and tilt determine how much of the terrain an operator is able to see. When a 16 mm lens (approximately 31 degrees horizontal field of view) was compared to a 4 mm lens (approximately 96 degree horizontal field of view) for teleoperation of a small vehicle in an indoor environment, there was a significant reduction in the number of times the vehicle touched obstacles with the 4 mm lens (Silverman, 1982). In a similar study by Horst, Rau, LeCocq, and Silverman (1983) using these same lenses, stereo viewing improved performance (number of obstacles bumped) for both camera lenses.

ACKNOWLEDGMENTS

This research was performed under Martin Marietta Astronautics Group IR&D task, "Planetary Rover Technology."

REFERENCES

Frederick, D.K., "An Autonomous Roving Vehicle For the Planet Mars," Technical Report MP-83, Rensselear Polytechnic Institute, Troy, NY, October 1984.

Horst, R.L., Rau, P.S., LeCocq, A.D., and Silverman, E.B., "Studies of Three-Dimensional Viewing in Teleoperated Systems," Remote Systems Technology Supplement to the Transactions of the American Nuclear Society, Detroit, Michigan, Summer Meeting, 1983.

Silverman, E.B., "Robotic Technology Experiments for Nuclear Power Plant Inspection and Maintenance," Transactions of the American Nuclear Society, Volume 43, 1982.

Spain, E.H., "Assessments of Maneuverability with the Teleoperated Vehicle (TOV)," Presented at the Fourteenth Annual Symposium of the Association for Unmanned Vehicle Systems, Washington, D.C., July 19-21, 1987.

Wilcox, B.H., Salo, R., Cooper, B., and Killon, R., "Computer-Aided Remote Driving," Thirteenth Annual Meeting of the Association for Unmanned Vehicle Systems, Boston, Massachusetts, July 21-23, 1986,.